Camera-Based Scene Reconstruction and Pose Estimation for Robotic Task Planning


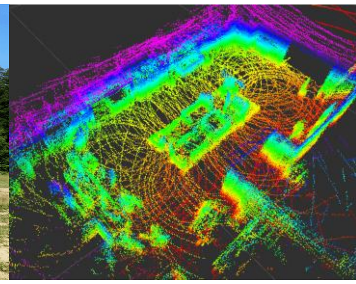


Company Background:

At Stratom, we are driving the future of automation by developing unmanned ground vehicles and autonomous robotic systems for commercial and defense applications — whether in safe, controlled settings or dynamic and challenging terrain.

Specializing in unmanned cargo movement, autonomous mobile robots (AMR) and robotic refueling, our proven tools, methods, technologies and strategic services continuously meet our customers' unique and evolving needs in logistics and operations. Our solutions enable them to reduce monotonous, difficult or dangerous tasks to optimize uptime and efficiencies, address labor shortages, increase profitability, and keep people safe.

Project Description:

Autonomous robots operating in the real world must build accurate models of their environment to plan and execute tasks reliably. While raw depth data from stereo cameras can capture the geometry of a scene, the resulting point clouds are often noisy, incomplete, and poorly suited for direct use in robotic task planning algorithms. Transforming this raw data into a clean, smooth, and semantically useful surface representation is a critical step toward enabling robots to reason about and interact with their environment.

This project challenges students to develop a perception pipeline that takes camera imagery as input and reconstructs the 3D geometry of a scene. Common targets, such as parked vehicles and other objects composed of mixed materials (e.g., glass, painted metal, rubber, and reflective surfaces), will be used to ground the problem in realistic operating conditions. Students will explore methods for point cloud and/or mesh reconstruction and registration, and develop refinement techniques to reduce noise, handle incomplete data, and produce representations that are usable for downstream tasks. Limited use of simulation environments (e.g., Gazebo) may be leveraged to generate data or evaluate approaches, but the primary focus is on real-world performance. The ultimate goal is to generate a reliable scene understanding that enables the estimation of surface poses for use alongside robotic scanning and planning operations.

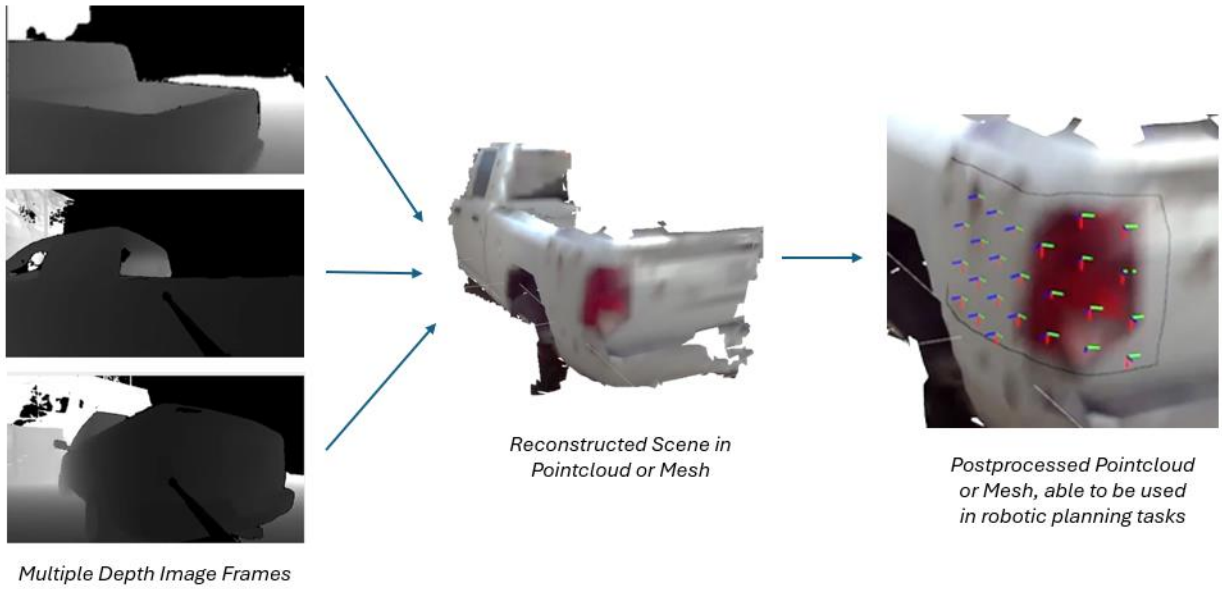

Figure 1: Example pipeline

Students will focus their efforts on:

- Collecting stereo image data using real stereo camera hardware across a variety of scenes and surface types
- Generating depth maps and converting them into 3D point clouds
- Registering and fusing data across multiple viewpoints to produce a consistent, unified scene representation
- Developing mesh refinement algorithms to smooth noisy geometry, fill holes, and reduce artifacts while preserving meaningful surface features
- Leveraging simulation tools (e.g., Gazebo) to generate synthetic data or validate reconstruction and refinement approaches, where appropriate
- Generating pose outputs (position and orientation) suitable for execution by a robotic manipulator
- Handling real-world challenges such as noise, occlusions, and variable lighting conditions
- Integrating the full pipeline into a ROS2 node for real-time processing and deploying it within a Docker-based development environment
- Visualizing input data, intermediate results, and final outputs using tools such as RViz

Desired Skillset:

- C++, Python
- ROS2
- Linux
- Docker
- 3D perception concepts (e.g., point clouds, reconstruction, registration, pose estimation)

This project is an excellent opportunity for students to build on their experience in C++ or Python while developing practical skills in 3D perception, scene reconstruction, and pose estimation. Students will gain hands-on experience working with real sensor data, integrating perception pipelines within ROS2, and producing outputs suitable for robotic interaction. While prior experience in all of these areas is not required, students should be prepared for fast-paced learning and working through real-world integration challenges.

Team Size: ~4 Students

Location:

Meetings will primarily be held remotely using Teams or Zoom. Our office is located in Louisville at 331 South 104th Street, Suite 235, where students may have the opportunity to get in-person support and/or and present a live demonstration of their results to the Stratom engineering team.

Note: All intellectual property developed as part of this project will be owned by Stratom